UC Berkeley
School of Information

Workshop Materials

Big Data: Values and Governance
March 31, 2014 – Berkeley, CA

Workshop Goals

The purpose of this workshop is to bring together a small but diverse group of people to grapple with complex value and governance issues that arise because of the "big data" phenomenon. You are being asked to consider both challenges and opportunities on the values side; and, with respect to governance, to think broadly about how values can be protected. We are asking you to roll up your sleeves, collaboratively work across perspectives, and brainstorm about these issues. You should focus on teasing out the tensions and offering a conceptual map of how to consider the values at stake, and then consider the solution space. When considering the solution space please note dependencies between various tools, for example, a technology might exist to protect a value, but its adoption might lag in the marketplace. What intervention might spur adoption? Publicity? Law? Collective action?

The more you can work across domains, the better. What matters in healthcare or education may be different than what's at play in national security or criminal justice. The challenges are different in public, private, and civil society sectors, even though the three are often intertwined. Whenever a path forward seems obvious, make sure to consider how it would play out in different domains or other sectors. Second-guess yourself, and others (gently). Be open to new ideas rather than the limits of existing debates. Step outside your comfort zone. Ideally, you'll identify emerging issues, not just those fully in play today.

There are no formal facilitators. Through the use of post-it notes and paper, and a short video, we hope you can capture your thinking, and produce a short summary of your conversation to share with the group during our closing session (See below). We'd appreciate your help in making sure that the conversation is both fruitful and documented.

Above all else, be respectful. You have been brought together because you have diverse and important perspectives, but this also means that there will be disagreement. Please attack this challenge with an open mind, a willingness to listen, and an appreciation for difference. The more that you can channel different voices and respect divergent perspectives, the more successful you will be in teasing out the values and governance issues that we need to address in an era of "big data." This will also advance our

conversation tomorrow. *Unlike tomorrow's public event, today we are under a version of Chatham House Rules—you are free to talk about what was discussed and we will be disclosing all participants, but you may not attribute remarks to another individual.*

The workshop backgrounders are intended to offer some meat for discussion. Each includes a description of the issues at play, a few short case studies, and a series of questions to help shape your conversation. The purpose of this document is to spark ideas and trigger conversations but you do not need to stay focused on this document in your conversation. Use it as a jumping off point to allow you to think about these issues from different angles.

Workshop Schedule

Big Data: Values and Governance
March 31, 2014 – Berkeley, CA

1-1:45 Plenary Session Room 210 South Hall
- Opening Remarks, Nicole Wong & Deirdre Mulligan
- Soapbox presentations by participants

1:45-2:00 Short break for transition

2:00-3:30 Workshops (food available) Rooms 107, 205, &210

3:30-3:45 Short break

3:45-4:15 Second Round Soapbox

4:15-5 Report Backs from Sessions and Wrap-up

UC Berkeley
School of Information

Workshop Materials 1
Predicting Human Behavior

Big Data: Values and Governance
April 1, 2014 – Berkeley, CA

Brief Description

Predictions can be made from trace data, with varying degrees of personal or societal consequence (e.g., biometric data used to predict fatigue and risk of accident, search engine queries used to predict hospital admission, gaming companies and casinos can predict compulsive gambling problems, government agencies can predict criminal activity). Big data allows us to see patterns that predict behavior with some level of accuracy, and use them, without understanding why. It can allow us to prevent harms or promote good policies without understanding how the pattern connected with them matters. But, it can also find patterns that segment and cluster in ways that re-encode socially and at times legally taboo discrimination along lines of race, gender, class, etc. The ease with which big data allows us to slice and dice the population for different treatment—from pricing, to border crossing, to advertisements—suggests that concerns are far broader than privacy, but instead reach individual meritocracy and fairness as well. The use of such classification to alter the information landscape—search returns, advertisements, political mailers, etc.—also raises specific questions about shared experiences and their potential constitutive role in community and democracy.

Detailed topic description
Large-scale data analytics offers the potential to connect data collection with previous behavioral research to detect patterns in how people behave, and predict outcomes based upon those patterns.

In commerce, such applications can be used to predict buying potential, attract new customers, prevent attrition, and target products, services, and offerings. It can also be used to predict individual behavior outside the commercial sphere. For example, the Internet radio service, Pandora, prides itself on customizing music playlists for listeners based on user feedback on music content. In order to customize music content, listener data is kept on an individual level. A listener may not realize that their song feedback may reveal other information about them: specifically political orientation. In previous election cycles, Pandora has used a political ad-targeting system to place candidate ads based on inferences from their previous music selections. Because users must provide a zip code when registering for Pandora service, inferring political orientation allows localized campaign ads to be placed. Pandora users might be surprised to know that their musical taste is used by campaigns to tailor outreach to them. Some have suggested that this sort of

targeting is useful in energizing voters. Recent campaigns have used technology and data to great effect. Others believe it is deeply problematic.

In the health sector analytics may be comprised of a variety of data including demographic information, personal habits, and genetic data. They may also take into account non-health data to infer health conditions or predict health outcomes. A patient's living arrangements or other aspects of their personal lives may be used as variables to determine how likely a patient is to be hospitalized in the near future, which offers healthcare providers an intervention opportunity resulting in saved lives and decreased cost. While this type of research offers valuable insights into optimizing health care, studies have shown patients may be de-anonymized with only a zip code, birthday, and gender. Beyond these inherent privacy risks, it is important to consider how these analytics may take away an individual's capacity to select what information is shared with whom.

The growing use of data brokers to obtain demographic and predictive information about people opens up further questions. Companies are using all sorts of non-traditional data to assess credit worthiness claiming they are a cheaper and unregulated alternative to credit reports. Regulators are exploring the privacy and consumer protection issues in this big data application. Consumers are largely unaware of these alternative scoring techniques and the impact they have on their experiences. Individuals have little to no access to the underlying data used to generate scores. These purchased scores could be used to determine the likelihood someone could pay their medical bills, thus determining someone's access to potential care based on metrics that may be incorrect.
Recruiters researching job candidates often search social media (LinkedIn, Twitter, etc.) and are able to discover attributes like race, age, gender, sexual orientation, or religious beliefs, which leaves open the potential for discrimination. Companies like Reppify are attempting to analyze and distill this information in order to assess job candidates, and offer a third-party option for companies to eliminate discrimination liability. Recently the FTC settled with Spokeo, Inc., a data broker that compiles and sells detailed information profiles on millions of consumers. The FTC had charged the company with violating the Fair Credit Reporting Act by failing to make sure that the information it sold would be used only for legally permissible purposes; failing to ensure the information was accurate; and failing to tell users of its consumer reports about their obligation under the FCRA, including the requirement to notify consumers if the user took an adverse action against the consumer based on information contained in the consumer report.

Predictive analytics may also be used to scan text to predict a person's suicide risk and present an opportunity for intervention. While this could undoubtedly save lives and improve counselors and healthcare providers in assisting patients, this could also cause harm if false positives were found using publically available information (i.e. social media). Non-health data, such as whether or not a person owns a cat or drives a minivan without having any children, is being used to target potential clinical trial participants. These predictions, used to locate and solicit participation in medical studies, are made without access to medical records.

Predictive variables from found data may be used to assess who is likely to do what. For example, Chicago's police department developed an algorithm inspired by sociological studies on criminal behavior that is designed to help locate who might be more likely to commit violent crimes. Those on the so-called "heat list" come under police supervision, whether or not they are under investigation for any past or present criminal activity.

In NYC as part of the Juvenile Robbery Intervention Program detectives monitor teenagers use of social media in an attempt to intervene before participants engage in criminal activity. The aim of the program is "to break the trajectory of those born into poverty and neglect, and winding up behind bars before their 18th birthday" and to reduce crime.

But big data isn't certain to yield useful or improved predictions. Google Flu Trend once hailed as an earlier indicator of the value of big data, has been found to yield less useful predictions than the CDC data traditionally used to make such predictions. In the context of national security, experts concluded that the collection and analysis of millions of telephone call detail records provided "little benefit over other intelligence activities." Importantly, they concluded, it did *not* directly contribute "to the discovery of a previously unknown terrorist plot or the disruption of a terrorist attack…in only one instance over the past seven years…arguably contributed to the identification of an unknown terrorism suspect." The researchers who critiqued Google Flu Trend warned that the "quantity of data does not mean that one can ignore foundational issues of measurement and construct validity and reliability and dependencies among data. The core challenge is that most big data that have received popular attention are not the output of instruments designed to produce valid and reliable data amenable for scientific analysis." These issues arise across big data applications, raising important questions about the fair and legitimate use of predictions.

Employers are also beginning to use predictive analytics to scrutinize employee working habits, work interactions, and assessing the likelihood of an employee to leave. Hewlett-Packard (HP) began calculating a "flight risk score" to predict how likely an employee is to leave their position. Data including salaries, raises, job ratings, and job rotations were analyzed historically and then used to create a predictive model that allowed HP to intervene with job rotations or raises.

Predictions come with the possibility of error. Misidentifying consumer buying habits may result in misplaced advertisements, but misidentifying health conditions could come with serious lifetime consequences, misidentifying the potential for criminal or terrorist activities even more so.

Drawing upon and combining different types of data to predict behavior and illuminate new insights into human behavior presents exciting potential as well as significant policy and legal challenges. Insights can be used in ways that are beneficial to society as well as deeply problematic.

Questions to consider

- What are the implications for systems of oversight and due process protections with regard to the use of predictive models (for example in employment, healthcare and policing)?
- How do predictions affect the criteria that figure in the police's "reasonable suspicion"?
- Where social network graphs, telephone calling records, location information or other transactional data that reveals potential associations among individuals are the basis of predictions what is the impact on freedom of association?
- How fit is available data for use in predicting terrorism? Crime? Are there aspects of these contexts that make big data approaches more or less problematic?
- How does speculative data play into current regulations on employer/employee relationships? Do these predictive analytics favor one demographic likely to exhibit certain behaviors (like frequently changing jobs) over another?
- With so many separate firms and organizations collecting data and forming predictions, is it reasonable to expect individuals be able to seek out and keep track of all conclusions being drawn about them – even if these results were open?
- What are the current types of predictive scores available to companies and how are they being used?
- How accurate are these scores and the underlying data used to create them?
- Are there new risks posed to those who are under-represented in the data? Are they ways to protect those excluded as well as included?
- Is the Data Quality principle in some versions of the Fair Information Practices a useful entry point for considerations of big data's predictive use? How might this concept be extended?
- What privacy concerns surround the use of predictive scoring? Would providing control over collection address them?
- What legal protections currently exist for consumers regarding the use of predictive scoring, both in the United States and internationally? Should some of these scores be considered eligibility determinations that should be scrutinized under the Fair Credit Reporting Act?
- Does the Equal Credit Opportunity Act provide a useful tool in dealing with predictive scoring?
- Are there particular consumer segments that are at risk in a market of prediction?
- What sorts of concerns arise from the ability to predict health status from non-health information? How do existing legal protections address them?
- Given that our understanding of what harms should be avoided, and what sorts of patterns are problematic, will evolve, what sorts of governance systems will best surface problems and aid in their resolution?
- Do our current statutory and Constitutional protections for privacy reflect the predictive power of big data?

UC Berkeley
School of Information

Workshop Materials 2:
Scrutinizing Algorithms

Big Data: Values and Governance
April 1, 2014 – Berkeley, CA

Brief Description
Concerns about the biases reflected in "automated data processing systems" and with the impact on human dignity of reducing individuals to prescribed data sets go back 40 years. Today, those same questions are being raised about algorithms. The concerns tend to center around the capacity of technical systems to 1) mask intentional abuse or bias; 2) embed the unacknowledged values of "technicians"; 3) skew legal/policy choices through simplification; 4) mask values through complexity.

Of course, bias and values embedded in systems are a product of human decision-making, and exist whether embedded in technology or not. The issue therefore might more productively be framed as one about: Whose biases? (a move to algorithms puts one set of actors in charge, laws and policies another) How explicitly they are stated? How transparent are biases to others (can they be interrogated)? Who can comprehend them? Is transparency sufficient? How easily can biases in a system be managed? What sorts of oversight and accountability mechanisms exist? How well do they scale? How can systems be tested? For example testers in anti-discrimination law, testing wrt software and hardware.

Accountability is fundamentally about checks and balances to power. In theory, both government and corporations are kept accountable through social, economic, and political mechanisms. Journalism and public advocates serve as an additional tool to hold powerful institutions and individuals accountable. But in a world of data and algorithms, accountability is often murky. Transparency is clearly useful, but it is not a panacea, as algorithms are complex, their interactions with data sets even more so, and even those few who can understand them are rarely aware of their full impact in practice.

Detailed topic description
A key motivation for the Privacy Act of 1974 was the growing use of automated systems, in the public and private sectors that processed citizens' data and made determinations about them. Suggestions from the 1973 Health, Education and Welfare report[1] included: separating data for statistical purposes from systems used to evaluate benefits, publishing data to allow for independent analysis, doing away with secret automated data processing systems, and provide individuals with access and correction rights to data in

[1] Department of Health Education, and Welfare. *Records, Computers and the Rights of Citizens*, 1973. https://epic.org/privacy/hew1973report/

administrative systems (those used to make decisions about individuals).
The opacity of algorithms and their effects is all the more apparent today, as more of our interactions — with governments, with our friends, with businesses — are computer-mediated and large-scale data analysis of those interactions and other behavior becomes more important.

The opacity of computer technology applied to important matters of public policy is, therefore, not new and transparency can butt up against other values, like intellectual property. For example, proprietary development of electronic voting machines has led to conflicts where state elections are run using algorithms that might not be transparent or open to inspection because of contract or trade secret law.

Questions to consider

- Beyond questions about whether the market is sufficient or governmental regulation is necessary, how should algorithms be held accountable?
- Do we need to avoid biases beyond those reflected in our legal commitments to equal protection? Can we articulate what biases systems (of humans and machines) ought to avoid?
- Are they ways to make sure biases are more explicitly stated?
- Are there persistent biases favoring or discriminating against certain groups? Are they affirmative things those designing or deploying algorithms could be asked to consider that might address this?
- Does transparency bolster comprehension? What other techniques might be used?
- What sorts of oversight and accountability mechanisms exist? How well do they scale?
- How can systems be tested? Are models used by regulators to detect intentional and disparate impact discrimination in areas such as housing and finance useful? Where an algorithm is not stable do these tools work? Do we need others?
- Privacy impact assessments have become a common tool of privacy professionals. They aim to bring privacy issues to the table early where they can influence data flows, market entry decisions, or interfaces. Can we imagine a tool that would assist us in evaluating algorithms? Who would use it? In an ideal world what sort of skills and training would they have?
- Should deviation from an algorithmic output be viewed as a sign of bias being introduced or mitigated? (Different regulators take different views on this now.) How can we tell which it is?
- What sorts of harms should algorithms be trained to avoid and how would one prove the harm was avoided? Can we imagine something akin to an immutable audit in this area?
- How do we square, or prioritize, competing claims to intellectual property and demands for accountability where proprietary algorithms are part and parcel of government operations?

UC Berkeley
School of Information

Workshop Materials 3:
Collection, Consent & Context

Big Data: Values and Governance
April 1, 2014 – Berkeley, CA

Brief Description

The growing quantity of networked devices and environments that generate data about the actions and interactions between people, their machines, and their environments is feeding the big data phenomena. From cars, to homes, to public streets and transportation, to private clubs, to MOOCs, data about individuals—where they are, who they are with, how fast they are moving—is on the rise and on the move, coursing through networks, analyzed in clouds, and largely invisible to those that generate the data. This information rich and networked environment presents opportunities for personal and collective good, however there are both conceptual and practical problems with privacy models centered on individual control exercised at the point of collection. Moreover, some of the public benefits of big data may be skewed by under inclusive or over inclusive data sets, raising questions about how data is collected, whom it is collected from and how providing privacy controls may impact outcomes. Moreover, this data has the potential to be used to the detriment of individuals and communities— regardless of whether it's used in a privacy sensitive manner.

Detailed topic description

As individuals go about their everyday lives, they accumulate an inordinate amount of data. Many transactions, communications, and interactions are traceable, allowing a patterned profile of human activity to form as a byproduct of people's activities. Instances of communication traces are connected to large corporations who provide the scaffolding and infrastructure for those communications to take place. Not only do individuals communicate via text, email, blogs, website commenting systems, and social media platforms, but they also engage in self-tracking behaviors, such as using fitness applications to monitor their daily routines.

In addition to these forms of data, telephone companies have records of people's locations, insurance companies know who received what medical service, and financial companies have records of purchasing patterns. Sometimes, the data collected is connected to personally identifiable information - including names, addresses, and phone numbers - and sometimes it is linked to less perfectly identifiable information, such as IP or MAC addresses, zip codes, or gender.

Marketers cobble together these bits of information to create a profile, but even innocuous seeming information like a zip code can be added to a person's birth date and gender in order to pinpoint an individual. According to Latanya Sweeney, a professor at Harvard, up to 87% of Americans are potentially identifiable from their zipcode, birthday, and gender. Geo-location may also be used to generate marketing data and generate personalized

content to consumers during their brick-and-mortar shopping experiences. Euclid Analytics is one example of a "plug and play" sensor that allows stores to track their consumers' location within 10 feet or as customers pass by the storefront, and understand their shopping habits at other Euclid-enabled stores. The consumer may not know this location information is being generated about them, or understand how their shopping habits are being used to tailor their in-store experience.

The growing use of "smart" appliances and gadgets, otherwise known as the Internet of Things, stands to expand the scope of personal data collection. Nest, a Wi-Fi enabled and "smart" thermostat and smoke detector, was recently acquired by Google and caused commotion as users worried about what the merger would mean for their personal data. Would Google be able to correlate their comings and goings in their house with their extensive databases including Gmail and Droid records?

License plate scanning has always been a part of the way police officers identify and patrol motorways. Now technology has enabled license plate scanners to automate this process and track cars and their drivers systematically at set or transient locations. These scanners are capable of recording the license plate number, some details about the car, and recording the time and location. This automation can assist police in finding stolen vehicles or locating abducted children, but it also creates a large database of scanned vehicles that were present at the scene of a suspected crime for use in an investigation. The data may also be of interest to private parties seeking to establish the where about of another individual.

Questions to consider

- What does it mean to protect privacy in this environment of ubiquitous collection?
- Are there lines that can be drawn between acceptable and unacceptable uses of data? Are they domain specific?
- How well do our existing policies and approaches respond to this Internet of things and the data it generates?
- What does transparency look like in this environment?
- What effect does the Internet of things have on data de-identification or anonymization?
- Do defaults in devices that emit data, and developer kits that allow access to data become more important in this environment?
- Could the concept of "privacy-by-design" provide specific guidance to hardware and software developers? Does this depend on whether they are general purpose or designed for specific contexts?
- What concerns arise from the potential capacity of government to continuously record activities occurring in public? Are existing legal frameworks able to address them? Are there concepts other than privacy that might be useful in placing limits on systematic government collection of otherwise publicly available information? Freedom of association?
- What responsibility does the average citizen have in curating their personal data trail? How should individuals be made aware of the ways in which their identity is

being linked to previously anonymous actions?

- What effects will surveillance have upon collective action and behavior in public places as public consciousness about new trends in data surveillance grows?

UC Berkeley
School of Information

Workshop Materials 4:
Inferences & Connections
The Surprising Implications of Other Peoples' Data

Big Data: Values and Governance
April 1, 2014 – Berkeley, CA

Brief Description
You can know someone by the company they keep. Big data and powerful analytics, in both small and large hands, are capable of inferring, with varying degrees of accuracy, information and relationships that individuals may not have chosen to reveal directly. From inferring sexual orientation or political affiliation through social network graph analysis, to relationships from genetic material, to affiliations through location information, other peoples' data can be mined to reveal your secrets. Algorithmic inferences are a cornerstone of the "big data" phenomenon. Privacy protections centered on individuals exercising control over their own information have little impact on the implications of inferences applied to them but drawn from other peoples' data.

Detailed topic description
There is much excitement surrounding the potential "big data" has for organizations to derive information such as demographic data, affiliations, or genetic predispositions based upon inferences or connections among open data. Private industry, researchers, and governments are presented with new opportunities to learn more about customers, subjects, and citizens not just at a group level, but also on an individual resolution. Collecting data, however, becomes problematic when those data were not intended to give away personal information. Inferences become particularly fraught with ethical issues when these inferences are drawn from other peoples' data (relatives, social connections, etc.) As these inferences become more intertwined with decision-making in the private and public sectors, what mechanisms (regulatory, professional, and organizational) should be used to guide the use of these inferences?

Social media platforms offer additional opportunities to infer ethnicity, sexual orientation, religious or political views, gender, age, or emotional state (i.e. depression, happiness, substance abuse). These personal traits and attributes may be derived from behavior on the web such as website browsing logs or "likes" on Facebook. Beyond an individual's behavior online, attributes like sexual orientation may be derived from placement within a social network of friends.

Geo-location data may be used to infer associations based on the time and place of collection, and may be used in conjunction with the geo-location of others to infer which individuals may be together, for how long, and how often. Social media platforms (e.g.

Rally, Twitter, or Foursquare) allow users to tag location and friends they are with – even without consent from those tagged.

Large complications of genomic data pose serious personal consequences for relatives – living, dead, or not-yet-born. The family of Henrietta Lacks, the woman whose cells were used without her consent to derive the HeLa cell line for medical research, recently came to an agreement with the National Institutes of Health (NIH). The NIH enacted a policy that gives the Lacks family some control over the full genomic sequence from HeLa cells after published medical research revealed the family had a predisposition to certain diseases. Recently, the Supreme Court ruled that police may collect the DNA of arrested individuals, and enter that information into a local or national database, even if that individual is not convicted for the crime they were picked up for. Given known racial imbalances in arrest rates, what implications does this collection of genetic data in law enforcement have not only for the individuals arrested but their families?

These data-oriented practices leave many questions open for debate. As our data continues to be collected, combined, and utilized in unforeseen (and often invisible) ways, how will individuals and researchers come to understand how these data should be used? Who should be accountable for flaws in these inference techniques? What implications do inferences drawn without the knowledge of the individual have upon society and the individual itself?

Questions to consider

- In many cases, inferences may be drawn by connecting individuals within a network of others, or by capturing data made public on a forum like Facebook for another purpose. Should these data be considered public? Is it ethical to use that data for research without the consent of the subjects?
- What rights do third parties have with regard to inferences drawn from information properly collected from someone else? What sorts of harm would those rights protect against?
- What rights, if any, do third parties have in the underlying data used to make inferences about them? Are there any contexts where rights should be extended? Or is it more profitable to think about prohibiting harmful uses? For example, what rights do relatives (living, dead and not-yet-born) have in the collection of DNA and other biological materials outside of a medical treatment context (e.g., at point of arrest, by consumer genetic informatics services, for scientific inquiry)?
- What sort of technical approaches might prevent privacy or other sorts of harms that potentially flow from big data inferences?
- What responsibilities do entities that collect data that imputes connections have to those who are implicated by association?
- Is the Data Quality principle in some versions of the Fair Information Practices a useful entry point for considering the obligations on those who use big data to make inferences about those not in the data set? How might this concept be extended?
- Is it ethical to use data generated by an individual that was not intended to divulge

private attributes (e.g. sexual orientation) for that purpose? What if the inference is incorrect? Should the individual of interest have the opportunity to know these inferences are drawn about them, or have the opportunity to correct this information?

UC Berkeley
School of Information

Workshop Materials 5:
New Knowledge, New Duties?

Big Data: Values and Governance
April 1, 2014 – Berkeley, CA

Brief Description

The mining of large data sets in health care, financial and educational sectors can identify individuals at risk. Predictive models may identify an increased likelihood of disease; information-seeking activities may reveal heightened risks to safety or financial security. The ability to derive new knowledge—insights data subjects may not be privy to—invites questions about whether and when this knowledge gives rise to new obligations to warn third-parties of risks, duties to re-contact research subjects, or changes in the scope of duties to notify patients or former clients given reduced transaction costs.

Detailed topic description

The predictive and inferential capabilities of large-scale data analysis are impressive and sometimes unexpected. Famously (or infamously), Target was able to determine the stage of pregnancy for a shopper when even her own family members were unaware, through statistical analysis not just of that single shopper's data, but behavioral and volunteered data of a large population of shoppers.[2] Nearby, the City of Oakland has pursued various programs for crime prevention that attempt to make use of data analysis and predictive inferences to reach out to potential criminals in advance or respond to incidents more promptly: sending "Dear John" letters based on crowd-sourced and database-checked collection of license plate numbers in neighborhoods with high incidence of prostitution;[3] an instance of the popular "Ceasefire" program for contacting potential gang members with resources;[4] and a controversial Domain Awareness Center for combining public and private data sources (security cameras, gunshot microphones, even online social media postings).[5] As new data is collected and inferential ability improves, researchers — in academia, law enforcement or in private industry — may find themselves learning about or predicting development of various health conditions, mental illness, financial distress or even involvement in gang violence. When "big data" allows researchers to make these predictions, what new obligations are created on the researchers themselves?

---

[2] Duhigg, Charles. "How Companies Learn Your Secrets." *The New York Times*, February 2012, sec. Magazine. http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html

[3] Walter, Shoshana. "'Dear John' Letters: New Tool to Fight Prostitution." *The Bay Citizen*, June 20, 2011. https://www.baycitizen.org/news/crime/dear-john-letters-new-tool-fight/

[4] "Ceasefire in Oakland." http://www2.oaklandnet.com/Government/o/Mayor/i/Ceasefire/index.htm

[5] "Framework for Discussion of a Policy for Privacy and Data Retention for the Joint City-Port Domain Awareness Center." February 11, 2014. http://www2.oaklandnet.com/oakca1/groups/cityadministrator/documents/memorandum/oak045549.pdf

Of particular interest may be those contexts where society and the law have established a duty of care. In healthcare, obligations to disclose information that might help patients is involved in the ongoing debate over ethics and the integration of research and treatment in a "learning healthcare system."[6] Regardless of whether data collection is an instance of research or of care, should *risk* be the standard for determining the ethical oversight needed?[7] As massive open online courses (MOOCs) become increasingly common, these same ethical questions arise for teachers. What obligation do teachers, courseware providers or the designers of MOOCs have to their students when data analytics reveal their struggles? Purdue University, for example, has found success in retaining students using "Signals" to communicate early feedback based on automated analysis.[8]

In the US academic setting, ethical review through the system of Institutional Review Boards has been required for conducting human-subjects research since the 1970s. The Belmont Report[9] was motivated by the infamous Tuskegee syphilis experiment, where scientists withheld not only penicillin treatment but also the diagnoses of the disease from participants in the study. Since 2011, the Department of Health and Human Services has been considering, and seeking comment on, changes to the Common Rule to account for changes in the research landscape, including privacy and security risks related to informatics and the prevalence of research outside of single university sites.[10]

While statistical analyses can allow for impressive inferences, data analysis can also be misleading or even harmful in individual cases. The Food and Drug Administration issued a warning over 23andMe's marketing of its genome extraction and analysis in part because the risks and benefits of that information to the consumer were unclear.[11] At the same time, a NEJM article praised one aspect of 23andMe's service: the ability for individuals to choose not to learn information about disease predispositions.[12] While data scientists may have an

---

[6] Institute of Medicine. "The Learning Health Care System in America." Accessed March 30, 2014. http://www.iom.edu/Activities/Quality/LearningHealthCare.aspx.

[7] Faden, Ruth R., Nancy E. Kass, Steven N. Goodman, Peter Pronovost, Sean Tunis, and Tom L. Beauchamp. "An Ethics Framework for a Learning Health Care System: A Departure from Traditional Research Ethics and Clinical Ethics." *Hastings Center Report* 43, no. s1 (January 2013): S16–S27. doi:10.1002/hast.134. http://onlinelibrary.wiley.com/doi/10.1002/hast.134/abstract.

[8] http://www.itap.purdue.edu/learning/tools/signals/, as cited in: Willis III, J E, J P Campbell, and M D Pistilli. "Ethics, Big Data, and Analytics: A Model for Application." *EDUCAUSE Review* 48, no. 3 (2013). http://www.educause.edu/ero/article/ethics-big-data-and-analytics-model-application.

[9] *The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research*. *Federal Register*. Vol. 44, 1979. doi:10.1089/blr.1993.12.868. http://www.hhs.gov/ohrp/humansubjects/guidance/belmont.html

[10] Advance Notice of Proposed Rulemaking for Revision to Common Rule: http://www.hhs.gov/ohrp/humansubjects/anprm2011page.html

[11] http://www.fda.gov/iceci/enforcementactions/warningletters/2013/ucm376296.htm

[12] Annas, George J, and Sherman Elias. "23andMe and the FDA." *New England Journal of Medicine* 370, no. 11 (2014): 985–988. doi:10.1056/NEJMp1316367.

obligation to disclose new information, must they also provide a right *not*-to-know? In some instances, predictive models may identify an increased likelihood of disease; in others information-seeking activities may reveal heightened risks to safety. The ability to derive new knowledge—insights data subjects are not privy to—invites questions about whether and in what instances this knowledge gives rise to new obligations with respect to duties to warn third-parties of risks, duties to re-contact research subjects, as well as potential changes in the scope of duties to notify patients or other former clients given reduced transaction costs. And what if the data subjects are not "research subjects" but rather customers or users?

Consider, a few examples: Researchers working with genetic data may incidentally discover that an individual carries a genetic variant creating increased susceptibility to a disease. Such incidental revelations about individual research subject's health status raise questions about the interaction between confidentiality guarantees and ethical or professional duties to limit suffering. Similarly, during a research study data may acquire new clinical significance due to work of those researchers or insights provided by others research. While the health profession and health researchers continue to develop guidelines and best practices to navigate these sometimes competing duties, these sort of incidental findings raise novel challenges—even in the context of a long tradition working with these competing values, big data is raising broad and at times novel challenges.

But outside traditional health research, incidental insights from big data present more novel conundrums. For example, consider research by Seth Stephensons-Davidowitz identifying evidence that victims of child abuse, and those that suspect it, may use Google search to seek information and assistance. Knowing that many cases of child abuse are unreported, what sort of action would we like Google or other search engines to take? Do they have an obligation to act? Should we create one? What would it look like? Information provision? Alert to an authority? How does this interact with protections and desires for privacy and confidentiality? Might it chill the information seeking activities of those at risk —decreasing their chances at securing help—or is it a singular opportunity for effective intervention?

Questions to consider

- What responsibilities do data collectors have to inform data subjects when they know more than the subjects know about themselves?
- Academic researchers have well-established guidelines for human subject research (e.g. IRBs). However, they are currently under review given changes in research. Are there protocols or frameworks that could guide the private sectors' collection of data for research? How might we learn from experiences with the common rule and IRBs?

UC Berkeley
School of Information

Workshop Materials 6:
Inequalities and Asymmetries

Big Data: Values and Governance
April 1, 2014 – Berkeley, CA

Brief Description
While law is sometimes used to address information asymmetries that interfere with a sound functioning marketplace, big data insights in the hands of a broad array of entities -- from merchants to educational institutions -- works a quantitative and qualitative change on how we think of symmetry. Concepts of undue influence, manipulation, preying on vulnerabilities exist in law, but they are viewed, generally, as the exception, the marginal case where *information really matters*. Big data has the potential to make profiling and prediction nearly ubiquitous and widely distributed, changing the background assumption (as false as it may be) of equal bargaining power between institutions and individuals. Even among institutions, the availability of data—technology, data scientists, etc.—is unevenly distributed. Some organizations, agencies, and sectors are better equipped -- financially, technically or politically -- to gather, use, and analyze data than others.

Questions to consider
- The federal government spurred the move to a smartgrid. This effort was part stimulus and part energy policy—the hope being that it would aid in energy conservation as well as better load balancing to avoid outages. Some of the public benefit of the smartgrid is dependent on use of the detailed data it generates. Privacy rules are surely important, but reaping the value of the data for energy conservation is likely to require entities other than the utilities to analyze and use it. What sorts of technologies and policies might support access in a privacy protective way?
- If big data is as powerful as some promised, what sorts of risks does its use by businesses in consumer transactions pose?
- In government to citizen interactions?
- What benefits are possible in each context?
- What sorts of techniques can be used to promote the use of data to curb reliance on stereotypes?
- What techniques might reduce the possibility of masking prejudices in big data patterns?
- If big data is transformative, what are the consequences of certain sectors of the economy having greater capacity to leverage it?
- Who benefits and who doesn't, both at the individual level and the institutional level?
- In what ways does data magnify or combat asymmetries in power? Are there areas where we should be particularly concerned about such asymmetries?

UC Berkeley
School of Information

Workshop Materials 7:
Private Life

Big Data: Values and Governance
April 1, 2014 – Berkeley, CA

Brief Description

It is increasingly difficult to draw traditional lines between public and private spaces, as distinctions between the government and the private sector, personal and work life erode from the standpoint of data collection and use. The movement of technologies between public and private places, work and home, and various social circles are also challenging individual's ability to successfully manage their privacy. The data generated by the Internet of things, social networking technology, and other private and commercial uses of technology are generating massive amounts of data that can be, and are, analyzed to reveal detailed portraits of individual's relationships, lifestyles, behavior, and beliefs. These troves of data collected by private entities can be made available to intelligence agencies, law enforcement, and even private litigants relatively cheaply and easily. Several members of the Supreme Court signaled concern with the state of 4th Amendment law given the ability of private sector location aware services to function as a tracking device. Similarly, many members of Congress, backed by civil liberties organizations and information and communications technology companies, are pushing for reforms to the statutory protections for electronic communications and stored records. The past several months of disclosures about U.S. intelligence programs revealed intelligence agency efforts to access corporate data—both through the-Foreign Intelligence Surveillance Court approvals and through hacking private networks.

Questions to consider

- Cars, homes are being wired bringing activities many conceive of as private into third-party hands, sometimes into public view. How does our current legal framework address this reality?
- Are proposed ECPA reforms sufficient for protecting individuals' reasonable expectations of privacy? Should we be expecting those expectations to evolve?
- The Internet and connected devices are blurring the lines between the commercial world—where purchases happen—and the private world—where uses of purchased goods and services occurs—and replacing it with a near-continuous feedback stream to the device/service providers that allows for their increased monitoring of post-purchase behavior. This monitoring has been the action of a recent FTC settlement, and a similar feature used by schools to locate computers came under fire. In that case, software on the machines had remotely captured over 30,000 photographs using the computers webcams and 27,000 screen shots.
    - Are there clear limits on what data companies may collect about activities in the home through consumer use of devices and services? Ought there to be?

If notice is provided, is that sufficient?
    - Is there adequate guidance on privacy to those who own devices, services, content that they provide to employers or lease or license to consumers to possess and interact with?
- Cars and on-board systems generate, collect, and transmit information about the activities of the driver and other occupants. This data has many beneficial uses, however a recent survey found that one third of the 2,039 respondents had privacy concerns about the technology in driverless cars, in particular who might be able to access and use the data collected, such as insurers. How is the auto industry responding? What laws have states adopted? Are the privacy approaches of the NHTSA vehicle-to-vehicle communication proposal useful in thinking through other privacy issues?
- Mobile devices bleed across spheres of public and private life. Are there particular risks posed by this blurring and exposure?
- The smartgrid will collect, and potentially expose, much more detailed knowledge of what individuals do in their homes to utilities and others who provide value-added services. The California Public Utilities Commission adopted a groundbreaking privacy and security rule for Smart Grid energy usage data. However, the rule does not address the privacy and security of data that consumers disclose directly to third parties or authorize third parties to obtain from the smart meters in the home. How should regulators, policymakers and consumer advocates approach the issues arising from this data collection?
    - Is this new comprehensive privacy and data security framework a useful model for other connected devices?